The Kellogg Company Optimizes Production, Inventory, and Distribution

GERALD BROWN
gbrown@nps.navy.mil

Operations Research Department
Naval Postgraduate School
Monterey, California 93943

JOSEPH KEEGAN
joe.keegan@kellogg.com

Kellogg Company, P.O. Box 3599
Battle Creek, Michigan 49017–3599

BRIAN VIGUS
brian.vigus@kellogg.com

Kellogg Company, P.O. Box 3423
Battle Creek, Michigan 49016–3423

KEVIN WOOD
kwood@nps.navy.mil

Operations Research Department
Naval Postgraduate School

For over a decade, the Kellogg Company has used its planning system (KPS), a large-scale, multiperiod linear program, to guide production and distribution decisions for its cereal and convenience foods business. An operational version of KPS, at a weekly level of detail, helps determine where products are produced and how finished products and in-process products are shipped between plants and distribution centers. A tactical version of KPS, at a monthly level of detail, helps to establish plant budgets and make capacity-expansion and consolidation decisions. Operational KPS reduced production, inventory, and distribution costs by an estimated $4.5 million in 1995. Tactical KPS recently guided a consolidation of production capacity with a projected savings of $35 to $40 million per year.

The Kellogg Company has been using a large-scale linear program, the Kellogg Planning System (KPS), for more than a decade to guide its operational (weekly), production, inventory, and distribution decisions for breakfast cereal and other foods. In addition, KPS helps Kellogg to make tactical decisions on budget-ing, capacity expansion, capacity reassign-ment, and other similar issues.

KPS models Kellogg's operations in the United States and Canada, with global operations under study. These operations include the production, inventory, and distribution of hundreds of items from Kellogg-owned and contracted plants out

INDUSTRIES—AGRICULTURE, FOOD
PROGRAMMING—LINEAR, APPLICATIONS

Report Documentation Page		Form Approved OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE **DEC 2001**	2. REPORT TYPE **N/A**	3. DATES COVERED **-**

4. TITLE AND SUBTITLE	5a. CONTRACT NUMBER
The Kellogg Company Optimizes Production, Inventory, and Distribution	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Naval Postgraduate School Department of Operations Research Monterey, CA 93943**	8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT **Approved for public release, distribution unlimited**

13. SUPPLEMENTARY NOTES

14. ABSTRACT

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT **SAR**	18. NUMBER OF PAGES **15**	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**			

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std Z39-18

to distribution centers (DCs) and to customers.

Many large companies like Kellogg employ some sort of enterprise resource planning system (ERP) to coordinate raw-material purchases, production, distribution, orders, and forecasted demand. Kellogg's ERP is largely a custom, in-house product, and KPS is a custom tool to complement that system. Models like KPS are also attractive within the commercially available ERPs of SAP, Oracle, JD Edwards, and others. Indeed, these ERP systems offer plug-in features for planning production, distribution, and inventory, for example, SAP's Advanced Planner and Optimizer [SAP 2001]. However, even these features may be inadequate [Hsiang 2001]. For instance, they may use rule-based heuristics to attempt to meet demand while ignoring capacity constraints and then iteratively refine the solution, using heuristics, to attempt to meet capacity constraints. These heuristics take costs into account in their rules but do not minimize costs or maximize profits.

In current vernacular, KPS is a point solution because it is tailored to solve problems for particular functional areas of the business. KPS uses optimization to find the best long-term, cost-minimizing, integrated production, inventory, and distribution plan—within the limits of modeling assumptions and data accuracy. ERPs account for the low-level influence of individual near-term transactions; in contrast, KPS is a high-fidelity, prescriptive model that is ideally suited to evaluating alternate systemwide scenarios.

The Kellogg Company is the largest cereal producer in the world and is a leading producer of convenience foods. In 1999, worldwide sales totaled nearly $7 billion. Kellogg began with a single product, Kellogg's Corn Flakes, in 1906 and developed a product line of well-known, ready-to-eat cereals over the years, including Kellogg's All-Bran (1916), Complete Bran Flakes (1923), Rice Krispies (1927), Variety Pak (1938), Raisin Bran (1942), and Corn Pops (1949). Kellogg continues to

Managers modify plans that don't quite fit the realities of the plant floor.

develop and market new cereals, but its recent thrust has been in convenience foods, best exemplified by Kellogg's Pop-Tarts and Nutri-Grain cereal bars. In addition, Kellogg has recently entered the health-food business. Kellogg produces hundreds of products that are sold as thousands of stock-keeping units (*skus*), and acceptable profit margins depend on producing these products and packaging these skus as efficiently as possible.

The Kellogg Company had long used spreadsheets and special software for materials requirements planning (MRP) and distribution resource planning (DRP). But by 1987 Kellogg realized that its expanding product line and geographically dispersed production facilities required some means of systematic, global coordination and optimization. KPS was the result. After a year of development, we installed prototypic software in 1989. Although KPS was intended primarily for operational planning, the process of initial testing inspired the first real applications, which were tactical. For example, Kellogg was

adding a new production facility to expand capacity and extend the product line, and it used KPS to compare the overall cost implications of locating that facility in one existing plant versus another.

We installed KPS in 1990 and developed it further over several years. Initially, we used only systemwide average costs for each plant, basing optimal solutions on differentials in transportation costs and on available production capacities, rather than on differentials in manufacturing costs at the various production sites. This helped smooth the transition from the then-current, decentralized production-planning process to a more centralized process guided by KPS. In particular, using true costs, KPS would have shifted overall production patterns dramatically, and actually carrying out this shift would have been impractical. By 1994, however, we had introduced true manufacturing costs into the operational model, and this produced savings of $4.5 million in 1995 [Scott 1999]. Currently, KPS is in use weekly for operational planning and almost daily for dealing with tactical issues.

Kellogg's Products and Operations

Kellogg operates five plants in the United States and Canada: Battle Creek, Michigan; Memphis, Tennessee; Omaha, Nebraska; Lancaster, Pennsylvania; and London, Ontario. It has seven core DCs in such areas as Los Angeles and Chicago, and roughly 15 co-packers that contract to produce or pack some of Kellogg's products. Customer demands are seen at the DCs and at four of the Kellogg plants. In the cereal business alone, the firm coordinates the production of about 80 products, while packaging, inventorying, and distributing over 600 skus at about 27 locations (plants, co-packers, or DCs) with a total of roughly 90 production lines (including coaters, puffing towers, and cookers) and 180 packaging lines. Optimizing this many decisions is clearly a formidable task. The production, inventory, and distribution activities behind getting a package of Kellogg's Variety Pak to market helps to illustrate.

Kellogg's Variety Pak, sku 05337, contains 10 small boxes of different cereals, for example Corn Flakes, Rice Krispies, and Froot Loops. The individual boxes (*individuals*) can be produced and packed at

Kellogg wants to ship fresh products.

Kellogg plants in Battle Creek, Omaha, Lancaster, or Memphis. But not all plants produce all individuals. All of the individuals at one plant compete for packaging capacity because they are packed in the same size box. Many of the basic production decisions are independent; for example, Mini-Wheats and Corn Pops share neither materials nor production facilities. On the other hand, Frosted Flakes are essentially coated Corn Flakes, and production constraints dictate that not all Corn Flakes in a production run can be coated. Thus, the production of these two products must be synchronized.

Variety Paks are assembled only in Battle Creek, so Kellogg must coordinate prior production, packaging, and shipping from the other three plants with the final assembly operations in Battle Creek. Each individual is an intermediate product or *constituent sku*, which may be inventoried

at its production location or at an assembly site, or it may be used to create an assembly directly upon production. To minimize costs, the company must account for differentials in production, packaging, inventory, and shipping costs. After assembling and packaging the Variety Paks into cases, the plant then ships them to one of the seven DCs.

This example does not illustrate all of the complexities of production and distribution at Kellogg. KPS must also take into account the following:

—Some skus are produced and packed by co-packers;

—A given product and package combination may be packed into several different case sizes, each yielding a different sku;

—Bulk product is sometimes produced at one location and shipped to another for packaging; and

—Constituents of certain products, for example, Mueslix, can be produced at one location and shipped in bulk to another location where they are processed together with other constituents to create a single product. This contrasts with an assembly of distinct products like the Variety Pak. Some constituents may come from Kellogg plants and some from co-packers. (KPS can model any number of levels of these intermediate products.)

The Basic Operational Linear Program

The operational version of KPS makes production, packaging, inventory, and distribution decisions at a weekly level of detail. The model is based on a linear-cost version of the production-planning model introduced by Modigliani and Hohn [1955] (the economic lot-sizing model of Wagner and Whitin [1958] is related), but uses the now-standard production, inventory, and demand recursion:

$$HOLD_t = HOLD_{t-1} + MAKE_t - Demand_t$$

for all time periods (weeks) t,

where $HOLD_t$ is the inventory for a single product at the end of period t (a decision variable), $MAKE_t$ is the production of the product during time period t (a decision variable) and $Demand_t$ is the exogenous demand for the product during period t (data) [Dantzig 1959; Zangwill 1969]. This recursion is embellished in KPS with multiple stages of production, multiple products and skus, multiple plants and DCs, shipping lanes between the plants and DCs, and various capacity constraints. Johnson and Montgomery [1974, Chapters 4.6–7] describe similar models. We assume all data are deterministic; inventory safety stocks (that is, deterministic lower bounds on inventories) help prepare for uncertain demands and unforeseen production problems. KPS does not model raw materials: We determined early in the model's development that the burden of maintaining data on raw materials would outweigh any improvements provided by their inclusion.

With some variations, we model each Kellogg plant or co-packer as a set of processing lines that produce products (for example, Corn Flakes, Rice Krispies, and Blueberry Pop-Tarts), which in turn feed a set of packaging lines that pack finished skus. An sku is defined by product, package size, and case size: For example, sku 00122 is a case of 12, 18-ounce packages of Corn Flakes. "001" is the product code for Corn Flakes, and "22" encodes package and case information. Finished skus are

placed in inventory or are used to meet demand assigned to the plant (Kellogg plants can act as their own DCs) or are shipped to another plant or DC. All customer demand is aggregated by sku and location, that is, plant or DC.

For each week of the 30-week planning horizon, the decision variables associated with a plant are the following:

$MAKE_{h,p,t}$—Production of products on a processing line and other facilities that form a production process, for instance, the klbs (pounds $\times 10^3$) of product 015 (Frosted Flakes) produced on processing

Kellogg runs KPS each Sunday morning to guide production decisions in week 2 and beyond.

line LL01 and frosted with sugar on coater LC01 at the p = Lancaster plant in week . The product 015 and two production facilities LL01 and LC01 form the production process h. Every product requires processing on one processing line but may also consume capacity on other facilities, such as a coater.

$PACK_{k,m,p,t}$—Packaging of skus on particular packaging lines, for example, the klbs of sku k = 00122 packed on packaging line m = LP09 at the p = Lancaster plant in week t.

$HOLD_{k,p,t}$—Inventories of skus, for example, the klbs of sku k = 00525 held in inventory at the p = Battle Creek plant at the end of week t.

$SHIP_{k,p,p',t}$—Shipments of skus to or from other plants and DCs, for example, the klbs of sku k = 00525 shipped from the p = Battle Creek plant, in week t, to

the p' = Los Angeles DC. For the most part, we model a shipment leaving in week t to arrive in week $t + 1$.

A DC may be viewed simply as a Kellogg plant with no production or packaging facilities, and thus a DC incorporates only inventory and shipping variables.

KPS currently models Kellogg's own plants and distribution centers and about 15 co-packers. Co-packers are non-Kellogg production facilities under contract to produce and package products designed by Kellogg and bearing the Kellogg label or to produce constituents for Kellogg products that undergo final processing at Kellogg plants. (Kellogg performs no co-packing for other companies.) A co-packer has no exogenous demand assigned to it.

Within a plant, basic constraints for each week require that the system:

(C1) Does not exceed processing line capacities;

(C2) Does not exceed packaging line capacities;

(C3) Packages all products produced in a week into skus during that week (these are flow-balance constraints between processing and packaging);

(C4) For each sku, balances inventory from the previous week plus current packaging plus incoming shipments with outgoing shipments, plus consumption in assemblies if this is a constituent sku, plus exogenous demand assigned to the plant;

(C5) Satisfies safety stock requirements with inventory of each sku at each plant;

(C6) Coordinates processing lines and packaging lines as needed during each time period. For example, we may require that the time spent packaging an sku does not exceed the time spent processing the

product from which that sku is derived.

With respect to these constraints, a DC is just a plant without production-related constraints and a co-packer is a just a Kellogg plant with no exogenous demand.

Constraints (C1), (C2), (C4), and (C5) are implemented as elastic goals that can be violated at a price: When an elastic goal constraint is violated, a linear penalty per unit of violation is assessed. The occurrence of such an elastic violation may indicate that a little overtime is needed, or it may signal a bottleneck that cannot be avoided. Either way, the model recommends a systemwide plan that is optimally adjusted to deal with all such problems over all locations and time periods.

KPS does not model raw materials but does model some intermediate products. An intermediate product is viewed as a constituent sku that can be shipped to other plants where it is further processed or combined and packed with other constituent skus to create a finished, assembled sku. The Variety Pak described earlier is one instance of a constituents-to-assembly recipe. Also, semiprocessed Rice Krispies, called bumped rice, are produced at one plant and shipped in bulk totes (labeled and modeled as an sku) to a co-packer to be further processed and packed into Rice Krispies Treats. KPS handles the packaging of assembled skus by straightforward modifications of the flow-balance and packaging-line constraints: The assembly and packaging of an assembled sku consumes only packaging capacity and draws constituent skus from inventory or concomitant packaging.

The basic, time-invariant data for KPS are:

—Product and sku codes, product-to-sku relationships, including recipes for assembled skus, and case weights;
—The identities of the processing and packaging lines at each plant, the products or skus that can be processed or packed on those lines, nominal yields in klbs per

For its size, KPS is curiously difficult to solve.

shift, nominal processing or packaging cost for each product or sku in dollars per klb;
—Inventory costs for each type of sku at each plant in dollars per case per week;
—Shipping costs (dollars per case) by lane; and
—Various per-unit penalties for unmet demand, unmet safety stock, and line overcapacitation.

Data that vary by week are:
—Production and processing line availabilities at each plant measured in shifts;
—Variations in nominal yields or costs to account for time-of-year effects (for example, it may take longer to dry certain products during humid summer months) or effects associated with new lines, products, or skus where yields typically improve for the first few weeks after commencing production;
—Estimated demands, in cases, for skus at Kellogg plants and DCs based on forecasts made by the marketing department; and
—Desired minimum inventory levels (safety stocks) at demand locations.

The basic objective of KPS is to minimize the total cost of meeting estimated demands. The full objective function in-

cludes penalty terms for violating processing and packaging capacities, for not meeting demand, and for not meeting safety stocks.

A fundamental assumption behind a linear program is that each decision variable may take on any value in a continuous range, but production and packaging decisions at Kellogg (and other manufacturers) are not that flexible. For instance, KPS might suggest that about one third of a shift of a low-demand sku be produced at some plant in each week of the planning horizon, but the plant manager requires a one-shift minimum for that sku because of setup overhead (production time lost because of required equipment adjustments). Theoretically, it is possible to add binary variables to KPS to handle such situations, but we have not yet done this because the model has been hard enough to solve as a simple linear program. (Technology is improving, however, and a mixed-integer version of KPS to handle production and packaging setups is on the drawing board.) Therefore, managers review KPS-suggested production plans to modify plans that don't quite fit the realities of the plant floor.

Meeting Uncertain Demand—Forecasts and Safety Stocks

Much uncertainty is associated with the data for a long-term production-inventory model, and for KPS, the greatest uncertainty is in actual demands for skus. In the first few weeks of a time horizon, demand numbers may be fairly accurate because they are largely based on firm orders from customers. But even at week three or four of the horizon, actual demands may depart substantially from the marketing de-

partment's original forecast. Nonetheless, the overarching goal of Kellogg, and thus KPS, is to meet these customer demands.

Perhaps KPS should be a multistage stochastic-programming model that directly handles uncertainty in demand, and possibly uncertainty in manufacturing yields and line availabilities. But such a model would require an unwieldy amount of data and would be too difficult to solve: As a deterministic model, KPS is large and can take several hours to run. A stochastic-programming version would require orders of magnitude longer to solve. So KPS simply uses planned safety stocks, that is, minimum inventory levels, as a buffer for uncertain demand. A huge body of literature addresses safety stocks in production-inventory models (for instance, Silver, Pyke, and Peterson [1998, Chapter 7] and the 103 references they list), but these models typically require strong probabilistic assumptions and do not extend to multistage, capacitated, production, inventory, and distribution models like KPS. KPS uses simple rules for setting safety stocks that have been tuned manually over time: Experience is a good teacher in this case.

In KPS, safety stocks for an sku are set only at locations that see demand for that sku. Nominally, the safety stock for sku s at location p in week t is the sum of demands there in weeks t and $t + 1$, or some other function of future demands. However, if an sku is to be promoted in a special advertising campaign starting in week t, the safety stock in week t is set as the sum of estimated demands in weeks t through $t + 4$, or some other horizon that is longer than that for an sku not being

promoted. This extra buffer is kept because the actual demand for a promoted sku is higher and more variable than for one that is not.

Safety stocks also attenuate undesirable end effects in KPS. In particular, a cost-minimizing, finite-horizon, production-inventory model will always try to drive inventories to zero at the end of the planning horizon. Even in a model with safety stocks, we do not trust a finite-horizon model's prognostications in the last few time periods; without safety stocks, the number of periods of untrustworthy results would be even greater.

The Rolling Horizon and Solution Persistence

Kellogg uses KPS in setting a rolling horizon [Schrage 1999, pp. 187–188]:

Multi-period models are usually used in a rolling or sliding format. In this format, the model is solved at the beginning of each period. The recommendations of the solution for the first period are implemented. As one period elapses and better data and forecasts become available, the model is slid forward one period. The period that had been number 2 becomes number 1, etc., and the whole process is repeated.

KPS has one difference, however: Production and packaging decisions in the first week are fixed, and it is largely the second week's decisions that are set in motion at the beginning of week 1. The main reason for this is that it takes time to get raw materials and packaging materials in place for production, but KPS does not model such materials. Thus, at the beginning of week 1, the production and packaging plan is locked in place, having been made the week before or earlier along with orders for any materials that may not have been on hand.

Fixing model variables is one method of enforcing solution persistence [Brown, Dell, and Wood 1997]. That is, we require that the solution of KPS covering, say, weeks t through $t + 29$, persist to some degree with respect to last week's solution covering weeks $t - 1$ through $t + 28$. If the time lag in ordering certain materials is longer than a week, variables may also be fixed in weeks beyond the first week of the horizon.

Fixing variables to given values is a strong form of persistence; KPS also uses less coercive techniques, such as requiring a variable to lie within a specific range or penalizing deviations of the variable from a target value. For instance, it is common for a plant manager, with guidance from KPS, to decide that his plant will pack a certain sku in week 3, say, of the planning horizon. However, he will let the model decide (for now) exactly how much of this sku to produce above a specified minimum level. We may also view safety-stock levels and penalties for not achieving them as a form of solution persistence. In general, KPS exploits persistence to (1) handle lead times of raw materials, (2) reduce volatility in suggested production and distribution plans as the model horizon rolls ahead, and (3) incorporate managerial knowledge into the production plan that is too complicated to model more explicitly.

The Tactical Linear Program

Even though we originally envisaged KPS as only an operational model, we have also developed a tactical version of KPS for long-range planning, on the order of 12 to 24 months. Kellogg uses long-range planning to develop plant budgets, investigate capacity-expansion issues, test

new DC locations for cost savings, and so on. The tactical model is identical to the operational one except that (1) time periods consist of four-week blocks called months, (2) transportation is typically treated as instantaneous, and (3) a special time-cascade solution technique helps deal with the limited product shelf lives.

Aggregating data and changing transportation delays is straightforward, but handling shelf lives is not. Kellogg wants to ship fresh products and, as a rule, products should reach customers (retailers) within four or five months of production

Estimated savings of $4.5 million per year accrued from following the model's recommendations.

so that they have plenty of shelf life remaining. Shelf life can essentially be ignored in a 30-week operational model, but it cannot be ignored in a 16-month tactical model: If solved as a monolith, a 16-month version of KPS could, conceivably, call for producing an sku in month 1 to meet a demand in month 16, and this would not be realistic.

Conceptually, it is not hard to model a production, inventory, and distribution system that tracks the age (or the use-before date) of inventory: If the useful life of a product is τ periods, create τ copies of the inventory balance constraints, inventory variables, and shipping variables, and index them by the vintage of the product they represent. Unfortunately, this would increase the size of tactical KPS nearly five-fold. Instead, we solve the standard model using a heuristic called a sliding

time window.

To implement the sliding window, we generate the standard model, solve months 1 through 5, fix the first month's variables, solve months 2 through 6, fix the second month's variables, solve months 3 through 7, and so on. In this way, the model solution cannot see demand beyond five months in the future and therefore will not try to produce any products meant for sale more than five months in the future. This is a heuristic, but users are convinced that it works well. An added benefit is that the tactical KPS is easier to solve than the operational version, and this is important when running a large number of what-if scenarios. (Brown, Dell, and Wood [1997] give more details.)

Operational KPS in Action

Planning personnel meet for a half day about six weeks prior to the start of a quarter to schedule production and packaging for that upcoming quarter. They will change the schedule produced many times as the start of the quarter gets closer and data estimates are revised. But having a long-range, visible target is important in managing the purchase of raw materials with long lead times, and for making adjustments to plant capacities to satisfy demand.

To prepare for the quarterly meeting, we solve a weekly model with a 30-week horizon using a starting point projected from the end of the current quarter's schedule. Planners then develop detailed, implementable schedules for each processing and packaging line using KPS's production and packaging quantities as targets. This is a manual effort aided by spreadsheets. For example,

—If KPS shows a processing line being heavily utilized for a particular product, planners will enforce a regular schedule with production every week on that line using a whole number of shifts;

—If KPS shows consistent, low levels of production on a line, they will aggregate production into a sequence of larger production runs, in a whole number of shifts, once every few weeks; and

—If KPS shows unmet demand for a particular sku, planners may schedule weekend overtime for production and packaging of that sku.

The scheduled production at each plant, by product and totaled over the quarter, usually conforms closely to that suggested by KPS. However, KPS might source a particular product at Battle Creek rather

---

When Kellogg completes this project, it estimates the savings will be between $35 and $40 million annually.

---

than Omaha because of a very small cost difference, a difference that planners realize is negligible. In this case, planners might schedule Omaha for that product to create a more flexible or balanced schedule for the plants or to make the schedule more closely follow budgetary guidelines established months before using tactical KPS.

As time passes, planners compare the schedule against a weekly run of KPS in order to adjust process quantities and timing and to identify any approaching risks of unmet demand; planners may look several months ahead. In the shorter horizon of four to six weeks, when they have little

flexibility to change the schedule, they review only large deviations from the KPS-suggested packaging plan for potential modification.

Kellogg runs KPS each Sunday morning (the end of week 0) to guide production decisions in week 2 and beyond. Most model variables for week 1 are fixed: By Sunday, it is too late to make changes to production plans for week 1 because raw materials and packaging materials for week 1 are already at the plant or on the way.

Data for the weekly run of KPS come from a variety of sources. Demand data comprise a combination of forecasts from the marketing department and firm orders. Structural data on line capacities, yields, and the like are averages compiled over time, with data on new lines taken from engineering estimates. New or overhauled lines will have start-up curves associated with them. That is, yields will improve for several weeks as operators gain experience with the lines and product changeovers become smoother. Additional data come from conferences with plant managers: For example, a packaging line, BP25, at Battle Creek unexpectedly may be scheduled for maintenance in weeks 3 through 5, or a run of at least 200 klbs of product 123 must be made in week 4 because raw materials are reaching age limits and must be processed, or the yield on processing line OL01 at Omaha must be reduced in weeks 2 through 4 because of projected humid weather.

A typical model has roughly 100,000 constraints, 700,000 variables, and 4 million nonzero coefficients. It is solved with the X-System [Brown and Olson 1994] in

two to four hours on a DEC Alpha computer with 512 megabytes of RAM or in less than 20 minutes on a 500 MHz Pentium III laptop. For its size, KPS is curiously difficult to solve, and we have done research to find out why. Kellogg's system has scant slack capacity; small changes in plans affect many facilities and time periods; and about 70 percent of the model's constraints are taut at optimality. These are tough linear programs.

One key to solving KPS efficiently is the X-System's generalized-network factorization [Brown and Olson 1994; McBride 1985]. In particular, a selection of up to 95 percent of constraints C1, C3, and C4 will have at most two nonzero elements per column and thereby form a large generalized-network submodel. Such a submodel is easily identified, and substantial computation savings accrue because an explicit basis inverse (or other explicit basis factorization) need not be maintained for that submodel.

Once KPS is solved, results are loaded into a database and checked for consistency. If data problems are revealed, they can be corrected and the model rerun before central planners receive the results on Monday or Tuesday.

It is difficult to quantify the savings Kellogg obtains by using KPS rather than earlier manual methods. However, when we first introduced KPS for weekly planning, management decided that production costs should be equalized across plants. That is, average production costs would be used at all plants so that no plant would be likely to have a severe increase or decrease in suggested production because of not-yet-verified production-cost

differentials. After planners became comfortable with KPS and verified data, they introduced actual costs into the model. At that time, 1994, it was estimated that savings of $4.5 million per year accrued from following the model's production, inventory, and distribution recommendations [Scott 1999].

Tactical KPS in Action

KPS is just as important for tactical planning as it is for operational planning. Some representative examples demonstrate its value in the tactical arena.

Prior to the start of each fiscal year, planners populate the KPS database with estimated plant-cost and throughput data and forecasted demands for the fiscal year plus six months. We run the model to determine the optimal sourcing of production to satisfy the forecasted demand for the fiscal year; the extra six months of data mediate undesirable end effects in the solution. The firm uses the information on

---

KPS has saved the Kellogg Company millions of dollars since the mid-1990s.

---

production volumes to then establish financial budgets within the plants, inventory space requirements within the DC network, and equipment projections for each transportation lane.

KPS plays an integral role in evaluating production capacity. By investigating the utilization of various processing lines, planners can identify opportunities for improvement. If they see that the utilization of the lines that make certain products is low, they may consolidate to reduce costs. Conversely, if a set of such lines is fully

utilized, they may seek additional capacity. Also, if a product is produced on multiple lines at different locations with some lines operating at less than capacity, but with the low-cost location operating at capacity, a capacity increase at that location may be justified. Managers must evaluate the potential savings in variable costs using KPS and compare them with the cost of the capital improvements.

A recent consolidation project exemplifies the use of KPS for capacity planning. A combination of declining sales and increasing yields of certain products on certain lines was leading to underutilization of other processing lines, and managers conjectured that they could make savings by closing down some of the under-utilized lines.

Multiple plants produced the relevant products, and no single line was completely idle. Initial runs of KPS also revealed that simply removing one of the lines with low utilization would be unwise, because the remaining lines would have too little capacity to fully support the business. So, we used KPS to explore sets of alternatives for shutting down a subset of the lines and increasing capacity on others. We undertook the study in two stages, first determining which lines to shut down and then deciding where to increase capacity. Ideally, we would look at both sets of decisions simultaneously, but the large number of combined options necessitated this ad hoc decomposition. We identified reasonable alternatives for shutting down lines by running KPS with data covering 18 months, with different combinations of lines removed and additional capacity spread across the open lines. We sent

these alternatives off for more detailed financial and engineering analyses. Those analyses, combined with an accounting of fixed and variable cost implications for the scenarios, led to a decision about which lines to close. Then managers had to decide how to increase capacity.

By using KPS, we determined the required overall capacity increase and gave that figure to Engineering. Engineering generated a list of implementable options that could deliver this increase and their costs. We then incorporated capacity information so that we could measure variable-cost impacts of the various scenarios. Combining this with the capital required for each option enabled managers to make a financial comparison and select the best option.

Having selected the consolidation plan to follow, we created a transition plan for implementing it. KPS helped us answer these questions:

—When could Kellogg take the lines targeted for capacity increases out of service and install new equipment? (Engineering provided information regarding the time required and start-up curve estimates upon completion.)

—When should production cease on the lines being eliminated?

—How much inventory should Kellogg build to support the business during the transition?

When Kellogg completes this project, it estimates the savings will be between $35 and $40 million annually.

Kellogg also uses KPS to determine where to produce new products, to assist regular capacity reviews, and to justify or avoid the manufacturing and distribution

cost impacts of various projects. For its North American cereal business, we run KPS about 30 times a month to answer these types of what-if questions.

Conclusions

After more than 10 years, KPS is still in development: Business never stops changing. Global operations will require some refinements and more flexible inputs. We will introduce binary variables to more accurately model the realities of line scheduling. We will more accurately model production and packaging operations that are tightly coupled; this should result in improved solution quality and possibly speed. We may also model some critical raw materials with long lead times.

In both its tactical and operational roles, KPS has saved the Kellogg Company millions of dollars since the mid-1990s. Kellogg is introducing KPS into Latin America to improve operations there, and it is studying a global model. The advent of the European Union has simplified cross-border operations in Europe, and the east coast of the United States and the coasts of Europe are getting closer all the time.

Acknowledgments

APPENDIX

The following linear program is a didactic specimen of the Kellogg Planning System. The constraints and variables correspond to a generic location. Kellogg plants exhibit all features described, distribution centers have no entities associated with production and packaging, and co-packers have no exogenous demands.

Indices

f—food (product).
k—stock-keeping unit (sku).
p—plant (or distribution center).
t—time period.
l—processing line (for foods).
m—packing line (for skus).
h—production process.
$f(h)$—the food produced by process h.
$H(l, p)$—processes h that use line l at plant p.
$K(m, p)$—skus k that are packed on line m at plant p.
$H(f, p)$—processes h that produce food f at plant p.
$K(f, p)$—skus k that are packed from food f at plant p.
$K'(k, p)$—skus k' that are assembled from (constituent) sku k at plant p.
$M(k, p)$—packing lines m that pack sku k at plant p.

Data and [Units]

α_{hlpt}—fractional shifts used on processing line l to produce one klb of food $f(h)$ at plant p during time period t [shifts/klb].
β_{kmpt}—fractional shifts used on packing line m to pack one klb of sku k at plant p during time period t [shifts/klb].
$\gamma_{kk'p}$—fractional klbs of sku k used to make one klb of sku k' [klbs/klb].
d_{kpt}—demand for sku k at plant p during time period t [klbs].
$\underline{hold}_{kpt}$—safety stock for sku k at plant p during time period t [klbs].
u_{lpt}—capacity of processing line l, plant p, time period t [shifts].
u'_{mpt}—capacity of packing line m, plant p, time period t [shifts].
$hold_{kp0}$—initial inventory of sku k, plant p [klbs].

Decision Variables

$MAKE_{hpt}$—klbs of product $f(h)$ produced using process h at plant p during time period t.
$PACK_{kmpt}$—klbs of sku k packed on line m at plant p during time period t.

$HOLD_{kpt}$—klbs of sku k in inventory at plant p at the end of time period t.

$SHIP_{kpp't}$—klbs of sku k shipped from plant p to plant p' at the beginning of time period t (nominally arriving in period $t + 1$).

Constraints

(C1) $\sum\limits_{h\in H(l,p)} \alpha_{klpt}MAKE_{hpt} \stackrel{\le}{=} u_{lpt}$ $\forall l, p, t.$

(C2) $\sum\limits_{k\in K(m,p)} \beta_{kmpt}PACK_{kmpt} \stackrel{\le}{=} u'_{mpt}$
$\forall m, p, t.$

(C3) $\sum\limits_{h\in H(f,p)} MAKE_{hpt} - \sum\limits_{k\in K(f,p)} \sum\limits_{m\in M(k,p)}$
$PACK_{kmpt} = 0$ $\forall f, p, t.$

(C4) $\sum\limits_{m\in M(k,p)} PACK_{kmpt}$
$- \sum\limits_{k'\in K'(k,p)} \sum\limits_{m\in M(k',p)} \gamma_{kk'p}PACK_{k'mpt}$
$+ HOLD_{kp,t-1} - HOLD_{kpt}$
$+ \sum\limits_{p'\neq p} SHIP_{kp'p,t-1} - \sum\limits_{p'\neq p} SHIP_{kpp't}$
$\stackrel{\cdot}{=} d_{kpt}$ $\forall k, p, t.$

(C5) $HOLD_{kpt} \stackrel{\ge}{=} \underline{hold}_{kpt}$ $\forall k, p, t.$

(C6) (…additional constraints omitted.)

(C7) $HOLD_{kp0} \equiv hold_{kp0}$ $\forall k, p, t.$

(C8) All variables nonnegative.

Objective

Minimize production costs + packing costs + inventory costs + shipping costs + penalties for processing line capacity violations, packaging line capacity violations, unmet demand, and unmet safety-stock requirements.

Constraints (C1) and (C2), respectively, constrain shifts of activity on each processing line and on each packaging line in each plant during each period. The relational operators "$\stackrel{\le}{=}$" and "$\stackrel{\cdot}{=}$" signify that each of these constraints is elastic: If an elastic constraint is violated, a linear penalty per shift of violation is assessed. Constraints (C3) balance production activities with packaging activities. Constraints (C4) accumulate net finished production to satisfy demand by sku, plant, and time pe-

riod. These constraints balance packaging, shipments received from other plants, and inventory from the previous period with consumption to create assembled skus, shipments to other plants, inventory going into the next period, and exogenous demand. Constraints (C5) specify elastic lower bounds as safety-stock levels for each sku at each plant during each period.

Constraints (C7) initialize inventory, and all variables are nonnegative (C8).

The objective is to minimize the total of all costs over the planning horizon, including any elastic penalties arising from violation of elastic goal constraints. Elastic penalties are incurred primarily for unmet demand but also sometimes when the user fixes certain packaging variables. In this latter case, line capacity may be violated.

References

Brown, G. G.; Dell, R. F.; and Wood, R. K. 1997, "Optimization and persistence," *Interfaces*, Vol. 27, No. 5, pp. 15–37.

Brown, G. G. and Olson, M. P. 1994, "Dynamic factorization in large-scale optimization," *Mathematical Programming*, Vol. 64, No. 1, pp. 17–51.

Dantzig, G. B. 1959, "On the status of multi-stage linear programming problems," *Management Science*, Vol. 6, No. 1, pp. 53–72.

Hsiang, T. 2001, "The illusion of power," *OR/MS TODAY*, Vol. 28, No. 1, pp. 34–36.

Johnson, L. A. and Montgomery, D. C. 1974, *Operations Research in Production Planning, Scheduling and Inventory Control*, John Wiley and Sons, New York.

McBride, R. D. 1985, "Solving embedded generalized network problems," *European Journal of Operational Research*, Vol. 21, No. 1, pp. 82–92.

Modigliani, F. and Hohn, F. E. 1955, "Production planning over time and the nature of expectations and the planning horizon," *Econometrica*, Vol. 23, No. 1, pp. 46–66.

SAP 2001, "SAP advanced planner and optimizer," retrieved from ⟨http://www.sap.com/apo/⟩, August 24.

Schrage, L. 1999, *Optimization Modeling with LINGO*, LINDO Systems Inc., Chicago, Illinois.

Scott, D. 1999, "Custom optimization models provide competitive advantage in the supply chain planning process," Paper presented at the Council of Logistics Management Annual Conference, Toronto, Ontario, Canada, October 17–20.

Silver, E. A.; Pyke, D. F.; and Peterson, R. 1998, *Inventory Management and Production Planning and Scheduling*, John Wiley and Sons, New York.

Wagner, H. M. and Whitin, T. M. 1958, "Dynamic version of the economic lot size model," *Management Science*, Vol. 5, No. 1, pp. 89–96.

Zangwill, W. I. 1969, "A backlogging model and multi-echelon model of a dynamic lot size production system—A network approach," *Management Science*, Vol. 15, No. 9, pp. 506–527.

Don J. Scott, Vice President, North America Logistics, Kellogg Company, One Kellogg Square, Battle Creek, Michigan 49016-3599, writes, "The development of the Kellogg Planning System (KPS) 12 years ago which is described in the paper entitled 'The Kellogg Company Optimizes Production, Inventory, and Distribution' was a major improvement to Kellogg's supply-chain planning capabilities. KPS replaced our existing MRP and DRP Systems as well as enhanced our long-range capacity planning capabilities for Kellogg's USA business. KPS's ability to guarantee optimized detail production and deployment plans by simultaneously considering cost and capacity for manufacturing, warehousing and deployment was a significant improvement over the traditional planning systems that it replaced. Early in the life of KPS, we were able to document multi-million dollar operating savings that were driven from improved week-by-week optimized manufacturing sourcing decisions.

"Over the years Kellogg and Insight have continued to enhance KPS's capability and expand its use across all business channels. KPS's planning capabilities have played a vital role in Kellogg's ability to continually reduce costs and inventory levels while improving service and capacity utilization throughout our total supply chain."